

What brings me here...

The Bootable Cluster CD in the Undergraduate Curriculum for High Performance Computing Education.

Paul Gray
Dept. of Computer Science

gray@cs.uni.edu

Work supported in part by NCSI, the
National Computational Science Foundation

Some of you might already know me...

- ❑ That means that my talk must somehow be different.
- ❑ But it also means that you might know my passion for teaching High Performance Computing
- ❑ Over the past FOUR years⁺, "we*" have been doing a road show of sorts, promoting the potential of reaching into the undergraduate curriculum with HPC.
- ❑ It may be starting to settle in finally.
 - × PITAC Report
 - × HPCWire report(*)
 - × HPCWire follow-up letters

The 2 challenges of teaching parallel computing

- You need a parallel machine
 - × Uniform
 - × Predictable
 - × Affordable
- You need something to run on the machine
 - × Computing, not IT
 - × Visual
 - × Run in time allotted
 - × Show speed-up

Lessons learned

- ❑ Students will reboot the machines
- ❑ There is no time when teaching for room/machine partitioning
- ❑ Students need to see speed-up speed up
- ❑ Scheduled speed up is not seeing speed-up
- ❑ System administrators do not always provide complete resources

BCCD cluster solution

- ❑ Runs in memory, doesn't touch hard drive
- ❑ Small footprint Linux
 - × X
 - × Standard network connectivity
 - × Standard parallel tools
 - × Debugging/profiling tools
 - × Network analysis tools

BCCD cluster solution

- ❑ Can set up local, secure DHCP and DNS servers
- ❑ X on every machine, constant load/network monitoring
- ❑ Every student can run jobs
 - × Good for learning, lousy for scheduling
- ❑ Check node status, recover from node failures
- ❑ Easily change cluster size, create sub clusters
- ❑ Students can take home what they have in class

Real science examples

- Span granularity of parallelism
 - × Embarrassingly parallel
 - × Domain decomposition
 - × Heavily coupled
- Balance time issues
 - × Run in class time w/ real visualization
 - × Show speedup
- Span range of computational approaches
 - × Monte Carlo
 - × Finite Difference
 - × Cellular Automata

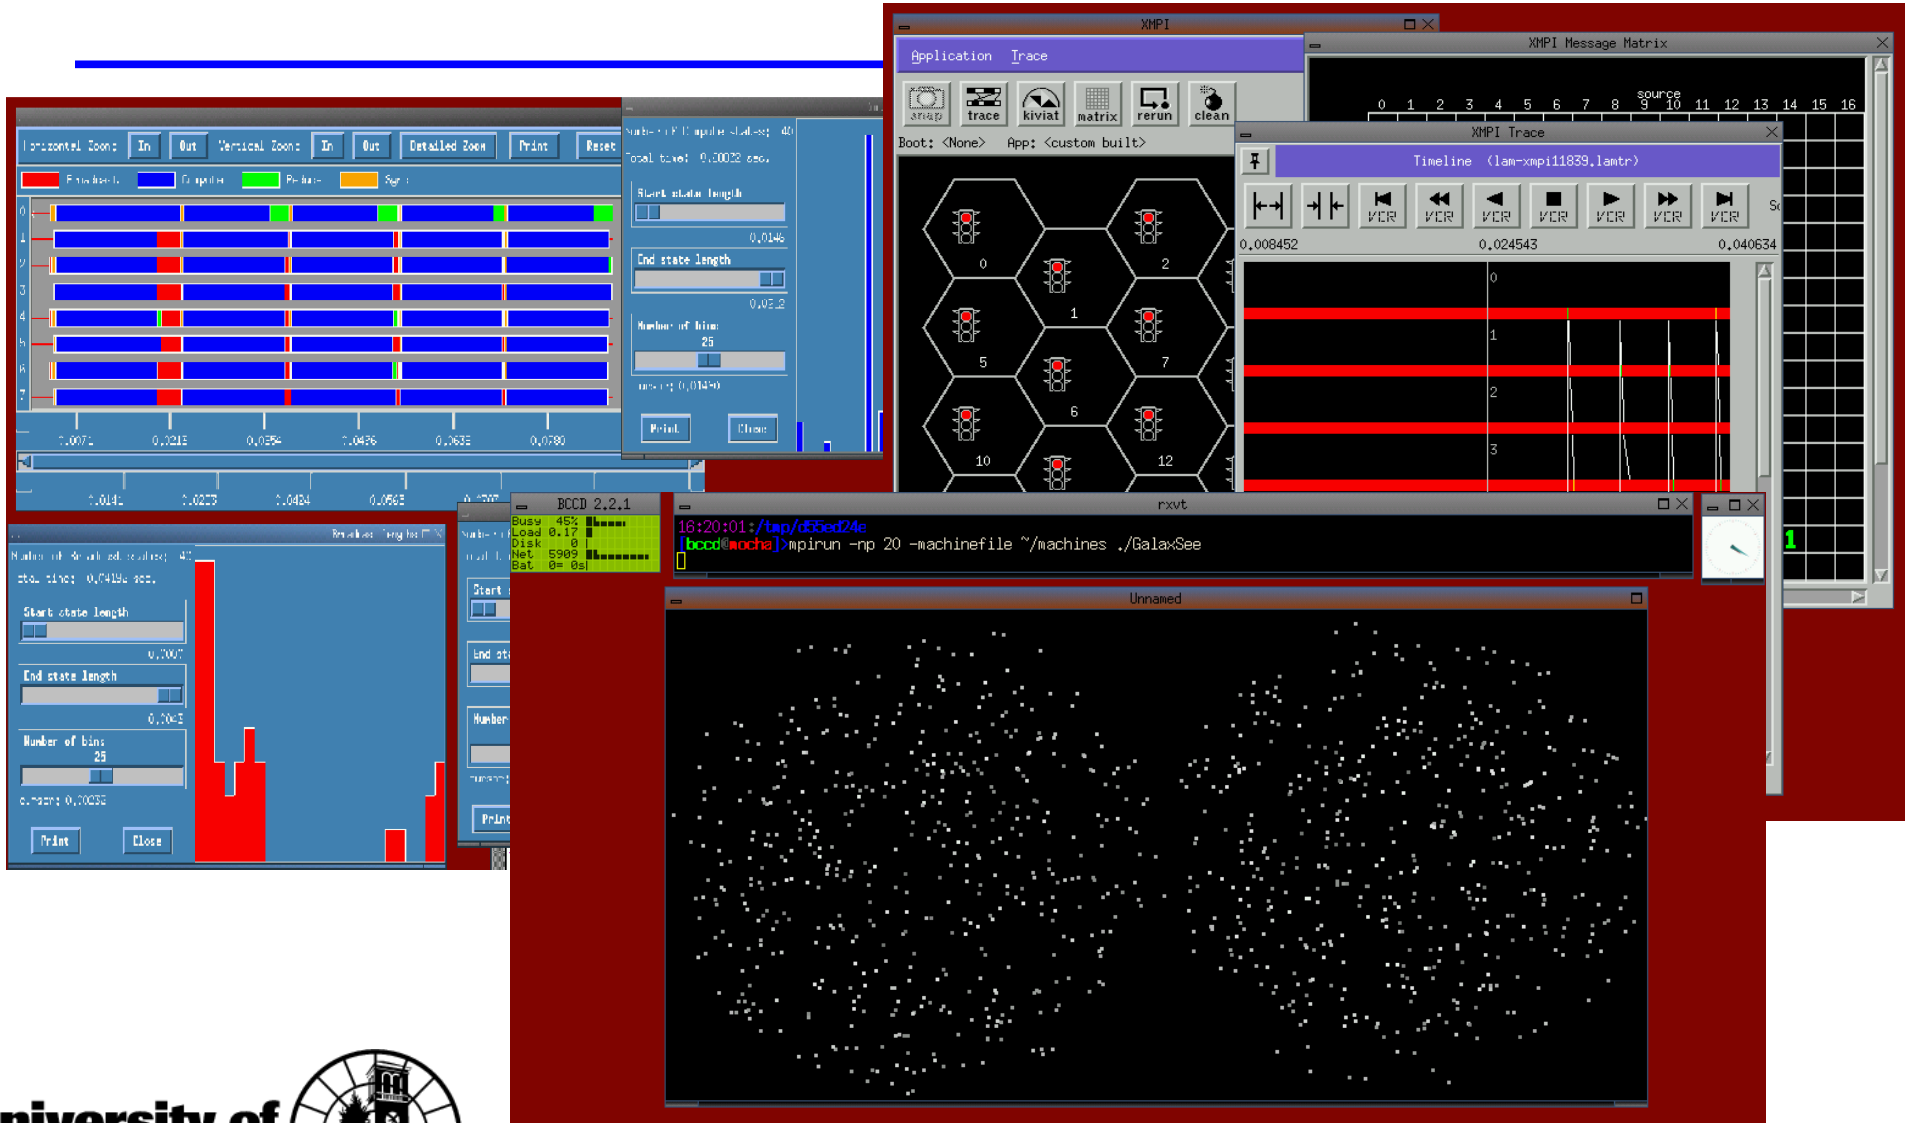
And if that's not enough...

- list-packages functionality:
 - × Customize your very own teaching materials
 - × Dramatic impact on “real HPC” look and feel
- Available everywhere! Students have access to (consistent) resources:
 - × In the classroom
 - × At home
 - × At the local Panera, Barnes & Noble, etc.

BCCD: The Bootable Cluster CD

- Wouldn't it be great if:
 - × You could **use existing computers** in the labs
 - × You had **no need to dual boot** or maintain multiple distros
 - × You could walk in to a computer lab with your class and have things "just right" for teaching HPC (in **9 CR's**)
 - × You had more than just the essential HPC tools, but also **visualization tools, debugging tools, profiling tools**, etc
 - × You didn't need to **patch, update, upgrade, test, fix, repeat**, ad nauseum to get stuff to work
 - × You had **no need to install the OS** from scratch
 - × You could **focus on curricular aspects** instead of administration
 - × the project that you heard about was **not vaporware**.

That is the BCCD



That's the BCCD

The image displays several overlapping windows from a system monitoring environment:

- XPVM (X P V M 1.2.5):** Shows a grid of Linux icons representing different coffee types (e.g., mexican, latte, housechi, guatemala, nacchia, italian, harrar). A trace event window is open, showing a list of events with filters.
- BCCD 2.2.1:** A terminal window showing system statistics:


```

      Busy 100%
      Load 1.79
      Disk 0
      Net 1419
      Bat 0= 0s
      
```
- openMosixview 1.5:** A network monitoring tool showing a table of nodes and a network diagram.

id	clusternodes	load-balancing efficiency	overall load	overall used memory	all memory	all cpu
all	all-nodes	96%	62%	10%	3436 MB	20
2571	192.168.10.11		!990			
2725	192.168.10.165		!990			
2726	192.168.10.166		!990			
2734	192.168.10.174		!990			
2730	192.168.10.170		!990			
2715	192.168.10.155		!990			
2716	192.168.10.156		!990			
2753	192.168.10.193		!990			
2737	192.168.10.177		!990			
2738	192.168.10.178		!990			
2739	192.168.10.179		!990			
2740	192.168.10.180		!990			
2741	192.168.10.181		!990			
2674	192.168.10.114		!990			
2700	192.168.10.140		!990			
2732	192.168.10.172		!990			
2728	192.168.10.168		!990			
2718	192.168.10.158		!990			
- Network Diagram:** A circular network topology with nodes labeled with IDs (e.g., 2700, 2732, 2728, 2718, 2747, 2751, 2725, 2726, 2734, 2716, 2715, 2730, 2753, 2716, 2737).
- Terminal Window:** Shows system messages and a command prompt:


```

      Warning: Permanently added 'bccd' (RSA) to the list of known hosts.
      Killed by signal 2.
      /mpich/bin/mpirun: line 1: 1249 Broken pipe          /home/bccd/Gal/./GalaxSee
      ee -p4pg /home/bccd/Gal/PI1168 -p4wd /home/bccd/Gal
      [bccd@host130:~/Gal]
      [bccd@host130] mpirun -np 2 -machinefile ../machines ./GalaxSee
      Warning: Permanently added 'host129.vmware.net' (RSA) to the list of known hosts
      [bccd@host130:~/Gal]
      [bccd@host130]
      
```

BCCD Features

- Built entirely from scratch, via web-fetched sources
 - × The build system is called "GAR"
 - × A mixture between LFS, Gentoo's emerge, and BSD's ports system
 - × GAR makes make make sense
- Big differences:
 - × *Everything* is cross-compiled to the target ARCH
 - × First step in the process is to build the toolchain compiler (i386-linxbccd-linux-gcc) and the libc subsystem for the TARGET.
 - × The entire system can be built with only user privileges

BCCD Features

- ❑ Totally non-invasive
 - × Access to local disk, but mounted ro by default*
- ❑ System overlay
 - × reboot and your system boots to the original OS
- ❑ Can be run on laptops, through vmware
- ❑ Thin client if there's a demand (from the academic side) or if there's funding in it.

BCCD Packages (small listing)

- ❑ PVM
- ❑ X-PVM
- ❑ MPICH (2.0 opt.)
- ❑ LAM
- ❑ XMPI
- ❑ C3 tools
- ❑ gcc, g++, gcj, g77
- ❑ mpicc, mpi77, mpiCC, mpic++
- ❑ hcc, hcp, hf77
- ❑ openMosix
- ❑ openMosixview
- ❑ openPBS (license?)
- ❑ electric fence, gdb, ddd
- ❑ vim, emacs, nedit, joe,
- ❑ vncserver, vncviewer
- ❑ ssh key sharing utilities
- ❑ stunnel and ipsec
- ❑ robotfindskitten
- ❑ ... many utilities

Future challenges

- Examples need to be based on current grand challenges, bench science
 - × Fewer contrived examples
 - × More community codes
- Examples need to span more of science
 - × Biomolecular and Bioinformatics
 - × Quantum chemistry and molecular modeling
 - × Physics
 - × Astronomy
- Examples need to be disseminated

Where can I find this cool stuff?

- ❑ BCCD
 - × <http://bccd.cd.uni.edu/>
- ❑ CSERD
 - × <http://cserd.nsd.l.org/>

Summary

- ❑ Those here would likely agree that HPC Education is needed at all levels, but especially undergraduate
- ❑ Resources to instruct on HPC topics are widely available, tailor to suit one's needs, environment and funding.
- ❑ The BCCD approach provides a flexible compromise between dedicated, standalone resources, administration headaches, and making use of current pools of resources.
- ❑ These slides, image downloads, and more info at:

<http://bccd.cs.uni.edu>